#13

Our Ref. No.: 042390P5379

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

RECEIVED
MAR 1 9 2002
Technology Center 2100

In re Application for:

Jerrold V. Hauck, et al.

Serial No.: 09/059,533

Filed: April 13, 1998

For: **EARLY ACKNOWLEDGMENT OF PRIMARY PACKETS**

Examiner: Torres, J.

Art Group: 2133

BEST AVAILABLE COPY

APPEAL BRIEF UNDER 37 C.F.R. § 1.192

Assistant Commissioner for Patents
Washington, D.C. 20231
ATTENTION: Board of Patent Appeals and Interferences

Dear Sir:

Applicants submit, in triplicate, the following Appeal Brief pursuant to 37 C.F.R. § 1.192 for consideration by the Board of Patent Appeals and Interferences (hereinafter referred to as "Board"). Applicants (hereinafter referred to as "Appellants") also submit herewith a check in the amount of $320.00 to cover the cost of filing this Brief. Please charge any additional amount due or credit any overpayment to deposit Account No. 02-2666.

TABLE OF CONTENTS

TABLE OF AUTHORITIES

I. REAL PARTY IN INTEREST

The real party in interest with regard to this appeal is Intel Corporation.

II. RELATED APPEALS AND INTERFERENCES

There are no other appeals or interferences known to the undersigned that will directly affect, be directly affected by, or have a bearing upon the Board's decision in the pending appeal.

III. STATUS OF CLAIMS

Claims 1-11 are pending in the application, all of which stand rejected. Claims 1-11 are on appeal.

IV. STATUS OF AMENDMENTS

There are no amendments pending.

V. SUMMARY OF THE INVENTION

The invention relates to bandwidth reclamation on a full duplex bus. A source node transmits a primary packet to a destination node. (Specification, p. 5, lines 1-14.) During the transmission, if the destination node is unable to receive the primary packet, the destination node will transmit a negative acknowledgement (NAK) packet to the source node. (See id.) Upon receiving the NAK, the source node will abort transmission of the primary packet. (See id.) Accordingly, the bandwidth that would have been used for transmitting the remainder of the primary packet can be reclaimed for another use. (See id.) Systems that employ acknowledgement gaps between the transmission of primary packets will also

benefit from this approach since the acknowledgement gap can commence as soon as the primary packet transmission is aborted. (Specification, p. 7, lines 11-16.)

Referring to Appellants' Figure 2a, source node 110 is transmitting PACKETA 116 to destination node 112. (Specification, Fig. 2a, p. 6, lines 25-27.) Destination node 112 cannot accept PACKETA 116. (See id.) This might be due to insufficient resources, for example. (Specification, p. 7, line 1.) Destination node 112 transmits a NAK 114 to source node 110. (Specification, p. 7, lines 4-16.) Upon sending the NAK 114, destination node 112 asserts an arbitration request 120. (See id.) In Figure 2b, source node 110 has received the NAK 114 and aborted PACKETA 116 before completing its transmission. (See id.) After aborting the transmission, source node 110 issues a bus grant 122 to destination node 112, thus allowing the bus time that would have been taken transmitting the remainder of PACKETA 116 to be used by destination node 112. (See id.)

Appellants' Figure 3 shows a system that includes a plurality of nodes 50-58. (Specification, Fig. 3, p. 7, lines 17-26.) Node 54 is transmitting a primary packet PACKETA out all of its ports. (Specification, Fig. 3, p. 8, lines 1-16.) The transmitting node, in this case node 54, is the nominal root node in the network topology and therefore receives bus requests from the other nodes such that it has the complete arbitration state of the network topology available to it. (See id.) Destination node 50 transmits a NAK to source node 54. (See id.) Upon receiving the NAK, node 54 aborts transmission of PACKETA and grants the bus to the highest priority requestor. (See id.) In this case, a bus grant by source node 54 would result in bus control going to either node 50, 52, 56, or 58. Again, the bus time that

would have been taken transmitting the remainder of PACKETA can be reclaimed by the node to whom the bus is granted.

VI. ISSUES PRESENTED

The issue presented in this Appeal is as follows:

(1) Whether Claims 1-11 are obvious over IEEE 1394 Standard for High Performance Serial Bus ("Reference 1") in view of U.S. Patent No. 6,185,184 to Mattaway, et al. ("Mattaway").

VII. GROUPING OF CLAIMS

Appellants assert that the claims do not stand and fall together. Rather, the claims are to be grouped as follows:

Group I:	Claim 1.	Group V:	Claim 6.
Group II:	Claims 2 and 3.	Group VI:	Claim 7.
Group III:	Claim 4.	Group VII:	Claims 9 and 10.
Group IV:	Claims 5 and 8.	Group VIII:	Claim 11.

Appellants will argue why each of these groups of claims should be allowed below.

VIII. ARGUMENT

A. Overview of the Invention and References

1) Distinctive Features of the Invention

The distinctive features of the invention include: (1) aborting a transmission without sending all of the primary packet; (2) generating a NAK concurrently while receiving a primary packet; (3) reclaiming bandwidth not used as a result of aborting transmission of the primary packet; and (4) granting the bus to a highest priority requesting node upon aborting the transmission. These features are shown in Figures 2-3 and described in the detailed description.

2) Overview of the Cited Art

Reference 1 discloses a standard for a high performance serial bus. (Reference 1, Abstract.) The serial bus architecture is defined in terms of communication protocols and nodes. (Reference 1, §3.1.) Nodes are logical entities with unique addresses. (See id.) The communication protocols are described as a hierarchy wherein a high-level transaction layer utilizes a low-level link layer. (Reference 1, Fig. 3-4, § 3.4.) The transaction layer defines high-level operations to read data from another node or write data to a node. (See id.) The link layer implements the transaction layer's read and write operations by providing a one-way 64 bit data packet transfer service with acknowledgement. (See id.)

Figure 3-11 in Reference 1 illustrates a typical packet transfer involving the transaction layer and the link layer. (Reference 1, Fig. 3-11, § 3.6.2.1.) A source node (or requestor) initiates a write request at the transaction layer which, in turn, causes the link layer to transmit a data packet to a destination (or responder) node. (See id.) Once the data packet is received in its entirety at the destination node link layer, the destination node link layer transmits an acknowledgement packet (ACK) to the source node. (See id.)

The process of transferring a packet via the link layer from a source node to a destination node is called a "subaction." (Reference 1, § 3.6.) A subaction may be asynchronous or isochronous. (See id.) Asynchronous subactions are comprised of a packet transmittal followed by an acknowledgement transmittal. (See id.) An isochronous subaction consists of a packet transmittal without an acknowledgment. (See id.) The acknowledgement packet contains a code to notify the source node that the destination node successfully received the packet or that the source node should resend the packet later. (Reference 1, § 3.6.2.4.) An acknowledgement code that indicates the destination node could not receive the packet is, for purposes of this discussion, a NAK. Thus, Reference 1 discloses that the link layer transmits a packet in its entirety and that an ACK/NAK is transmitted to the source node thereafter.

Mattaway discloses a protocol for establishing point-to-point communications between WebPhone users over a computer network. (Mattaway, Abstract.) A WebPhone application programming interface (API) enables WebPhones to communicate with each other. (Mattaway, col. 17, lines 11-35.) The WebPhone API utilizes sockets to communicate with other processes on a network. (See id.) A

socket provides a transmission conduit that can utilize a number of underlying communication protocols, including IP, UDP, RTP, and TCP. (See id.) These underlying protocols handle the details of transmitting and receiving packets. (W. Richard Stevens, *TCP/IP Illustrated,* Vol. I, 1994, pp. 2-3.) In other words, higher level application programs that use sockets may implement their own application-specific protocols (e.g., the WebPhone API) without having to worry about the underlying details of packet transmission and reception.

An application-specific protocol disclosed in Mattaway defines four high-level packets: <INFO REQ> , <INFO ACK>, <INFO> and <INFO ABORT>. (Mattaway, Fig. 17b, col. 26, lines 5-35.) A WebPhone stores one or more of a first name, last name, company, city, state, and country values in a Query field contained within the <INFO REQ> packet. (See id.) The WebPhone API opens a socket and sends the <INFO REQ> packet to a server as illustrated by message 1 of FIG. 17B. (See id.) After receiving the packet, the server extracts the values specified in the Query field of the <INFO REQ> packet and transmits an <INFO ACK> packet back to the WebPhone as illustrated by message 2 of FIG. 17B. (See id.) The server then automatically transmits one or more <INFO> packets to the WebPhone, as illustrated by messages 3A-C of FIG. 17B. (See id.) At any time following transmission of the <INFO ACK> packet, the WebPhone may transmit an <INFO ABORT> packet to either prevent transmission of any <INFO> packets or to stop transmission of any remaining packets, as illustrated by message 4 of FIG. 17B. (See id.)

B. Errors of Law and Fact

The Examiner failed to establish a *prima facie* case of obviousness in view of the references of record. The Federal Circuit Court of Appeals in In re Rijckaert, 9 F.3d 1531, 28 U.S.P.Q. 2d 1955 (Fed. Cir. 1993) held that:

> In rejecting claims under 35 U.S.C. §103, the examiner bears the initial burden of presenting a *prima facie* case of obviousness. . . . "A *prima facie* case of obviousness is established when the teaching from the prior art itself would appear to have suggested the claimed subject matter to a person of ordinary skill in the art." . . . If the examiner fails to establish a *prima facie* case, the rejection is improper and will be overturned. (Emphasis added.)

9 F.3d at 1532, 28 U.S.P.Q. 2d at 1956.

The Examiner employs inappropriate hindsight reconstruction guided solely by Appellants' disclosure. It is well established that the prior art must be viewed without reading into that art the Appellants' teachings. In re Sponnoble, 160 U.S.P.Q. 237, 243 (C.C.P.A. 1969). "The issue is whether the teaching of the prior art would, in and of themselves, without the benefit of appellant's disclosure, make the invention as a whole obvious." As the supplied references fail to teach or suggest numerous elements contained within Appellants' claims, it is only by inappropriate hindsight that one might possibly combine these references in the first instance.

Examiner has inappropriately combined Reference 1 and Mattaway as there is no suggestion of the desirability of the combination of those references even when their disclosures are viewed in their entirety. Specifically, Reference 1 discloses a low-level serial bus architecture (transaction, link, and physical layers) while the relied upon portions of Mattaway disclose a high-level means for obtaining

WebPhone client data from an information server. Accordingly, one of ordinary skill in the art would not have been motivated to combine these unrelated references.

The Examiner asserts that the cited references in combination teach or suggest aborting a transmission without sending all of the primary packet. The Examiner asserts that since the primary packets of Reference 1 are comprised of *quadlets* (groups of 4 bytes), *if* a primary packet could be aborted during transmission, it would occur on a quadlet boundary. (Office Action dated 11/7/01, p. 2, paragraphs 3-4.) The Examiner then relies on Mattaway to show that a WebPhone e-mail message is composed of multiple <INFO> packets and that its transmission can be aborted on an <INFO> packet boundary. (Mattaway, Fig. 17b, col. 26, lines 29-35.) The Examiner concludes that it would have been obvious to modify Reference 1 by including the abort feature from Mattaway such that the transmission of a primary packet in Reference 1 could be terminated on a quadlet boundary rather than sending all of the primary packet. However, this conclusion is erroneous.

First, Mattaway and Reference 1 do not teach or suggest partially transmitting a packet. Mattaway discloses that an e-mail message can be aborted on an <INFO> packet boundary. (Mattaway, col. 26, lines 30-33.) And although Reference 1 discloses that a primary packet is composed of quadlets, as discussed above, the link layer of Reference 1 only transmits packets in their entirety (i.e. all 64 bits). Furthermore, Reference 1 does not disclose a means for aborting transmission of a packet. Hence, neither Mattaway nor Reference 1 teach or suggest aborting the transmission without sending all of the primary packet.

Second, Mattaway and Reference 1 cannot be combined to achieve this functionality. According to Reference 1, quadlets have no header information whereas the <INFO> packets of Mattaway have a session ID field. (Reference 1, § 2.2.69; Mattaway, col. 33, Tables 7-8.) The <INFO ABORT> packet of Mattaway uses the session ID field contained in the <INFO> packet to identify which session to abort. (Mattaway, col. 19, lines 60-67.) Since quadlets have no identification information, the abort mechanism of Mattaway would not work. Furthermore, as discussed above, Reference 1 does not disclose the capability to abort a partially transmitted packet. Hence, there would be no motivation to combine the teachings of Mattaway and Reference 1.

Finally, the Examiner also asserts that the cited references in combination teach or suggest sending a NAK to the originator of the primary packet concurrently with the receiving. As discussed above, Mattaway discloses that the WebPhone API utilizes sockets to communicate with other processes on a network. (Mattaway, col. 17, lines 11-35.) As such, underlying protocols (e.g., IP, UDP, TCP, etc.) handle the details of transmitting and receiving packets while the WebPhone API implements the high-level WebPhone protocol. Hence, Mattaway cannot teach or suggest sending a NAK concurrently with receiving a primary packet since the primary packet is received in its entirety by the underlying protocol before the application program using the WebPhone API ever sees it. Likewise, Reference 1 teaches that asynchronous subactions are comprised of a packet transmittal followed by an acknowledgement transmittal. (Reference 1, § 3.6.) Thus, the cited references in

combination fail to teach or suggest sending a NAK to the originator of the primary packet concurrently with the receiving.

1) **Specific limitations of Group I, Group V, and Group VIII not described in the prior art**

All three groups require aborting transmission of a primary packet.

2) **Explanation why such limitations render Group I, Group V, and Group VIII unobvious over the prior art**

All three groups contain the same basis for independent patentabilty. But since each group depends from a different base claim, each group stands or falls individually. No reference cited by the Examiner teaches or suggests partially transmitting a packet. As argued above, Mattaway and Reference 1 do not teach or suggest partially transmitting a packet. Mattaway discloses that an e-mail message can be aborted on an <INFO> packet boundary. (Mattaway, col. 26, lines 30-33.) Reference 1 discloses that the link layer only transmits packets in their entirety (i.e. all 64 bits). Furthermore, the relied upon portions of Reference 1 does not disclose a means for aborting transmission of a packet. Hence, neither Mattaway nor Reference 1 teach or suggest aborts a transmission of the primary packet when a NAK is received. Therefore, the Examiner has failed to establish a *prima facie* case of obviousness and the Board should overturn this rejection.

3) <u>Specific limitations of Group II not described in the prior art</u>

Throughout the following argument, it is assumed that dependent claims carry with them the arguments made in favor of base claims and any intervening claim.

Claim 2 requires reclaiming bandwidth not used as a result of aborting transmission of a primary packet.

4) <u>Explanation why such limitations render Group II unobvious over the prior art</u>

As argued above, the relied upon portions of <u>Reference 1</u> and <u>Mattaway</u> do not teach or suggest aborting transmission of a primary packet, therefore it follows that neither teach or suggest reclaiming bandwidth as a result of aborting. Even assuming for the sake of argument that <u>Mattaway</u> could be construed as suggesting aborting transmission of a primary packet (which it does not), there is no suggestion of reclaiming bandwidth. There are many reasons why a transaction may be aborted (e.g., to save processor cycles, etc.), thus aborting the transaction does not necessarily mean reclamation of bandwidth. Therefore, the Examiner has failed to establish a *prima facie* case of obviousness and the Board should overturn this rejection.

5) <u>Specific limitations of Group III and Group VII not described in the prior art</u>

Throughout the following argument, it is assumed that dependent claims carry with them the arguments made in favor of base claims and any intervening claim.

All three groups require generating a NAK concurrently while receiving a primary packet.

6) **Explanation why such limitations render Group III and Group VII unobvious over the prior art**

Both groups contain the same basis for independent patentabilty. But since each group depends from a different base claim, each group stands or falls individually. As discussed above, Mattaway discloses that the WebPhone API utilizes sockets to communicate with other processes on a network. (Mattaway, col. 17, lines 11-35.) As such, underlying protocols such as IP and UDP handle the details of transmitting and receiving packets while the WebPhone API implements the high-level WebPhone protocol. Therefore, Mattaway cannot teach or suggest generating a NAK concurrently while receiving a primary packet since the primary packet is received in its entirety by the underlying protocol before the application program using the WebPhone API ever sees it. Reference 1 teaches that asynchronous subactions are comprised of a packet transmittal followed by an acknowledgement transmittal. (Reference 1, § 3.6.) Hence, Reference 1 also fails to disclose sending a NAK concurrently with receiving a primary packet. Therefore, neither Mattaway nor Reference 1 teach or suggest this limitation. Thus, the Examiner has failed to establish a *prima facie* case of obviousness and the Board should overturn this rejection.

7) <u>**Specific limitations of Group VI not described in the prior art**</u>

Claim 7 requires granting the bus to the highest priority requesting node upon aborting the transmission.

8) <u>**Explanation why such limitations render Group VI unobvious over the prior art**</u>

As argued above, <u>Reference 1</u> fails to teach or suggest aborting transmission of a packet. Thus, it follows that <u>Reference 1</u> fails to teach or suggest granting a bus upon aborting a packet transmission. In addition, <u>Mattaway</u> does not disclose a bus arbitration scheme. Therefore, neither <u>Reference 1</u> nor <u>Mattaway</u> teach or suggest that a <u>source node grants the bus to a highest priority requesting node upon aborting the transmission</u>. Therefore, the Examiner has failed to establish a *prima facie* case of obviousness and the Board should overturn this rejection.

IX. CONCLUSION AND RELIEF

Based on the foregoing, Appellants request that the Board overturn the Examiner's rejection of all pending claims and hold that all of the claims of the present application are allowable.

Respectfully submitted,

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP

Dated: 3/7/02

By: ______________________

Daniel J. Burns, Reg. No. 50,222

12400 Wilshire Boulevard
Seventh Floor
Los Angeles, California 90025
(310) 207-3800

CERTIFICATE OF MAILING:

I hereby certify that this correspondence is being deposited as First Class Mail with the United States Postal Service in an envelope addressed to: Assistant Commissioner for Patents, Washington, D.C. 20231 on 3/7/02.

Diane Martinez — 3/7/02 *Date*

APPENDIX A

The claims involved in this Appeal are as follows:

IN THE CLAIMS

1. A method comprising:

transmitting a primary packet from a source node towards a destination node on a full duplex bus;

receiving a NAK while the primary packet is being transmitted; and

aborting the transmission without sending all of the primary packet.

2. The method of Claim 1 further comprising:

reclaiming bandwidth not used as a result of aborting.

3. The method of Claim 2 wherein reclaiming comprises:

granting the bus to a highest priority requesting node; and

beginning transmission of a next primary packet from the highest priority requesting node.

4. A method comprising:

receiving a primary packet at a destination node;

identifying, during the receiving, that the node cannot successfully accept the primary packet; and

sending a NAK to the originator of the primary packet concurrently with the receiving.

5. A system comprising:

a full duplex bus;

a source node coupled to the bus, the source node to transmit a primary packet; and

a destination node coupled to the bus, to receive the primary packet, the destination node to generate a NAK if the primary packet cannot be successfully accepted, the NAK generated concurrently with the receipt of the primary packet.

6. The system of claim 5 wherein the source node aborts a transmission responsive to the NAK.

7. The system of claim 6 further comprising:

a plurality of additional nodes coupled to the bus to form a tree topology wherein the source node grants the bus to a highest priority requesting node upon aborting the transmission.

8. The system of claim 5 wherein an inability to accept the primary packet is caused by unavailability of a needed resource.

9. An apparatus comprising:

a transceiver;

a state machine coupled to the transceiver, the state machine to generate NAK in response to an inability to successfully accept a primary packet, the NAK generated concurrently with an ongoing arrival of the primary packet.

10. The apparatus of claim 9 wherein the inability to accept is caused by resource unavailability.

11. The apparatus of claim 9 wherein when the apparatus is a source of a primary packet, it aborts a transmission of the primary packet when a NAK is received.

TCP/IP Illustrated Volume 1

The P

W. F tevens

ADDISON-WESLEY PROFESSIONAL COMPUTING SERIES

The publisher offers discounts on this book when ordered in quantity for special sales.
For more information please contact:

Corporate & Professional Publishing Group
Addison-Wesley Publishing Company
One Jacob Way
Reading, Massachusetts 01867

Library of Congress Cataloging-in-Publication Data
Stevens, W. Richard

TCP/IP Illustrated: the protocols/W. Richard Stevens.
p. cm. — (Addison-Wesley professional computing series)
Includes bibliographical references and index.
ISBN 0-201-63346-9 (v. 1)
1. TPC/IP (Computer network protocol) I. Title. II. Series.

TK5105.55S74 1994 — 93-40000
004.6'2—dc20 — CIP

ISBN 0-201-63346-9
Text printed on recycled paper.
1 2 3 4 5 6 7 8 9 10 CRW 9796959493
First printing December 1993

Application	Telnet, FTP, e-mail, etc.
Transport	TCP, UDP
Network	IP, ICMP, IGMP
Link	device driver and interface card

Figure 1.1 The four layers of the TCP/IP protocol suite.

Each layer has a different responsibility.

1. The *link* layer, sometimes called the *data-link* layer or *network interface* layer, normally includes the device driver in the operating system and the corresponding network interface card in the computer. Together they handle all the hardware details of physically interfacing with the cable (or whatever type of media is being used).

2. The *network* layer (sometimes called the *internet* layer) handles the movement of packets around the network. Routing of packets, for example, takes place here. IP (Internet Protocol), ICMP (Internet Control Message Protocol), and IGMP (Internet Group Management Protocol) provide the network layer in the TCP/IP protocol suite.

3. The *transport* layer provides a flow of data between two hosts, for the application layer above. In the TCP/IP protocol suite there are two vastly different transport protocols: TCP (Transmission Control Protocol) and UDP (User Datagram Protocol).

   TCP provides a reliable flow of data between two hosts. It is concerned with things such as dividing the data passed to it from the application into appropriately sized chunks for the network layer below, acknowledging received packets, setting timeouts to make certain the other end acknowledges packets that are sent, and so on. Because this reliable flow of data is provided by the transport layer, the application layer can ignore all these details.

   UDP, on the other hand, provides a much simpler service to the application layer. It just sends packets of data called *datagrams* from one host to the other, but there is no guarantee that the datagrams reach the other end. Any desired reliability must be added by the application layer.

   There is a use for each type of transport protocol, which we'll see when we look at the different applications that use TCP and UDP.

4. The *application* layer handles the details of the particular application. There are many common TCP/IP applications that almost every implementation provides:

   - Telnet for remote login,
   - FTP, the File Transfer Protocol,
   - SMTP, the Simple Mail Transfer protocol, for electronic mail,
   - SNMP, the Simple Network Management Protocol,

   and many more, some of which we cover in later chapters.

If we have two hosts on a local area network (LAN) such as an Ethernet, both running FTP, Figure 1.2 shows the protocols involved.

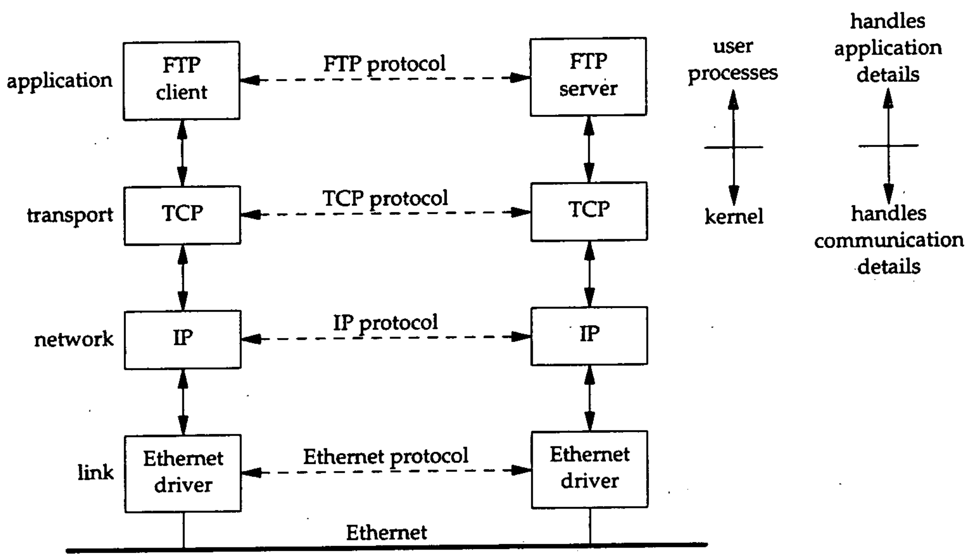

Figure 1.2 Two hosts on a LAN running FTP.

We have labeled one application box the FTP *client* and the other the FTP *server*. Most network applications are designed so that one end is the client and the other side the server. The server provides some type of service to clients, in this case access to files on the server host. In the remote login application, Telnet, the service provided to the client is the ability to login to the server's host.

Each layer has one or more protocols for communicating with its *peer* at the same layer. One protocol, for example, allows the two TCP layers to communicate, and another protocol lets the two IP layers communicate.

On the right side of Figure 1.2 we have noted that normally the application layer is a user process while the lower three layers are usually implemented in the kernel (the operating system). Although this isn't a requirement, it's typical and this is the way it's done under Unix.